

**Assessment Title**  
**Scenario-Based Activity Design and Implementation**  
**(AWS)**

CCC602 | Assessment-4

**The Student Name | Geni Bahar**  
Student ID: 27085022

**Tutor/Accessor | Ovesh Vohra**

Programme  
**Diploma in Cloud Computing and Cyber Security (120 Credits)**

Course  
**CCC602: Data Engineering in the Public Cloud**  
(Level 6, 30 Credits)

**Date: 12<sup>th</sup> February 2026**

## Table of Contents

Task 1: Cloud Infrastructure Setup and Automation for Data Ingestion, Cataloguing, and Real-Time Notifications Using AWS .....	2
Activities: 1.1: IAM and S3 Setup .....	2
Activities: 1.2: AWS Glue Configuration .....	7
Activities 1.3: Real-Time Notification Using AWS Lambda and SNS .....	19
Task 2: Query and Visualize the Dataset Using Athena and QuickSight.....	27
Activities: 2.1: Amazon Athena Queries .....	27
Activities: 2.2: Amazon QuickSight Dashboards .....	35
Task 3: Multi-Cloud Integration (GitHub → AWS → Azure Synapse Analytics).....	41
Activity 3.1: Data Export from AWS S3 .....	41
Activities 3.2: Upload CSV files from Azure Portal.....	45
Activities 3.3: Ingest Data into Synapse.....	46
Activities 3.4: Run SQL Queries in Synapse.....	48
Task 4: Final Report on AWS Services and Architecture.....	52
Activities 4.1: Report writing .....	52
Introduction .....	52
Overview of AWS Services Used .....	52
How These Services Work Together (End-to-End Workflow).....	53
Automation and No-Code Features.....	53
Data Engineering Perspective .....	53
Athena vs Azure Synapse SQL – Brief Comparison.....	54
Challenges and Solutions .....	54
Summary and Learning Outcome Reflection.....	54

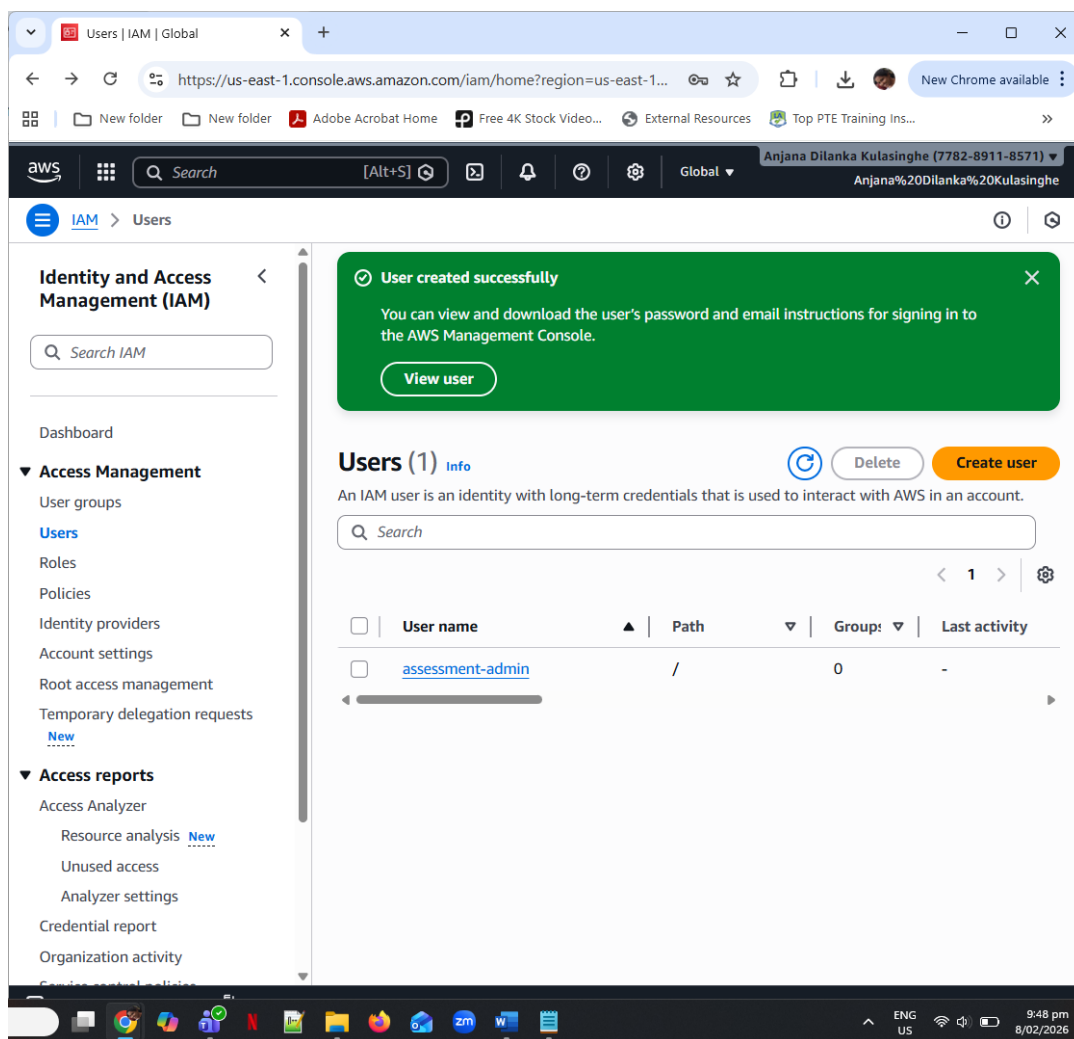
# Task 1: Cloud Infrastructure Setup and Automation for Data Ingestion, Cataloguing, and Real-Time Notifications Using AWS

In this task, I configured the core AWS services required to build the foundation of the automated data pipeline. The main objective was to securely upload datasets into Amazon S3, catalogue them using AWS Glue, and implement real-time monitoring using Lambda and SNS. This step established the storage, security, and automation components required for later querying and visualisation.

## Activities: 1.1: IAM and S3 Setup

An IAM user was created with AdministratorAccess permissions for the purpose of this assessment environment. Although in a real enterprise scenario, least-privilege access would be applied, full access is simplified configuration during the learning process.

An S3 bucket was created and structured using partition-style folders such as /orders/snapshot\_day=2017-01-01. This structure is important because AWS Glue detects partitions based on folder naming conventions. The January 2017 CSV dataset was uploaded into the correct snapshot folder to simulate raw data ingestion.



The screenshot shows the AWS IAM console interface. The browser address bar indicates the URL: `https://us-east-1.console.aws.amazon.com/iam/home?region=us-east-1...`. The user is logged in as 'Anjana Dilanka Kulasinghe (7782-8911-8571)'. The page title is 'assessment-admin' under the 'Users' section of IAM.

**Identity and Access Management (IAM)**

Search IAM

Dashboard

**Access Management**

- User groups
- Users**
- Roles
- Policies
- Identity providers
- Account settings
- Root access management
- Temporary delegation requests New

**Access reports**

- Access Analyzer
  - Resource analysis New
  - Unused access
  - Analyzer settings
- Credential report
- Organization activity

**User Details:**

- ARN: `arn:aws:iam::778289118571:user/assessment-admin`
- Console access: Disabled
- Access key 1: [Create access key](#)
- Created: February 08, 2026, 21:48 (UTC+13:00)
- Last console sign-in: -

**Permissions** | Groups | Tags | Security credentials | Last Accessed

**Permissions policies (1)** Remove Add permissions

Permissions are defined by policies attached to the user directly or through groups.

Filter by Type:  All types

<input type="checkbox"/>	Policy name <a href="#">↗</a>	Type	Attached via <a href="#">↗</a>
<input type="checkbox"/>	<a href="#">AdministratorAccess</a>	AWS managed - job f...	Directly

**Permissions boundary (not set)**

**Generate policy based on CloudTrail events**

CloudShell | Feedback | Console Mobile App | Privacy | Terms | Cookie preferences

© 2026, Amazon Web Services, Inc. or its affiliates.

The screenshot shows the AWS S3 console interface. At the top, a green notification banner states: "Successfully created bucket 'yoobee-assessment-data'. To upload files and folders, or to configure additional bucket settings, choose View details." Below this, the "General purpose buckets" section is active, showing a table with one bucket:

Name	AWS Region	Creation date
<a href="#">yoobee-assessment-data</a>	US East (N. Virginia) us-east-1	February 8, 2026, 21:52:08 (UTC+13:00)

Below the table, there are two summary cards: "Account snapshot" and "External access summary", both with "Updated daily" labels and "View dashboard" buttons. The browser's address bar shows the URL: https://us-east-1.console.aws.amazon.com/s3/buckets?region=us-east-1. The system tray at the bottom indicates the time is 9:52 pm on 8/02/2026.

The screenshot shows the AWS S3 console interface. At the top, a green notification bar states: "Successfully created folder 'snapshot\_day=2017-01-01'". The breadcrumb navigation shows the path: Amazon S3 > Buckets > yoobee-assessment-data > orders/. Below the notification, the folder name "orders/" is displayed, along with a "Copy S3 URI" button. The "Objects" tab is selected, showing a list of objects. The list contains one entry: a folder named "snapshot\_day=2017-01-01/". Above the list, there are action buttons: "Create folder", "Upload", "Copy S3 URI", "Copy URL", "Download", "Open", "Delete", and "Actions". A search bar is present with the placeholder text "Find objects by prefix". The system tray at the bottom shows the time as 9:53 pm on 8/02/2026.

The screenshot shows the AWS S3 console's 'Upload' interface. The browser address bar indicates the URL: `https://us-east-1.console.aws.amazon.com/s3/upload/yoobee-assessment-dat...`. The breadcrumb navigation shows the path: `Amazon S3 > ... > yoobee-assessment-data > orders/ > snapshot_day=2017-01-01/ > Upload`. The user is identified as 'Anjana Dilanka Kulasinghe' with a phone number '(7782-8911-8571)'. The main heading is 'Upload Info'. Below it, instructions state: 'Add the files and folders you want to upload to S3. To upload a file larger than 160GB, use the AWS CLI, AWS SDKs or Amazon S3 REST API. [Learn more](#)'. A dashed box contains the text: 'Drag and drop files and folders you want to upload here, or choose **Add files** or **Add folder**.' Below this, a section titled 'Files and folders (1 total, 35.1 KB)' contains a table of items to be uploaded. The table has columns for Name, Folder, Type, and Size. One file is listed: 'orders\_jan\_2017.csv' with a size of 35.1 KB. Below the table, the 'Destination' section shows the path: `s3://yoobee-assessment-data/orders/snapshot_day=2017-01-01/`. Other sections include 'Destination details', 'Permissions', and 'Properties'. The Windows taskbar at the bottom shows the time as 9:57 pm on 8/02/2026.

**Upload** Info

Add the files and folders you want to upload to S3. To upload a file larger than 160GB, use the AWS CLI, AWS SDKs or Amazon S3 REST API. [Learn more](#)

Drag and drop files and folders you want to upload here, or choose **Add files** or **Add folder**.

**Files and folders** (1 total, 35.1 KB) Remove Add files Add folder

All files and folders in this table will be uploaded.

< 1 >

<input type="checkbox"/>	Name	Folder	Type	Size
<input type="checkbox"/>	orders_jan_2017.csv	-	text/csv	35.1 KB

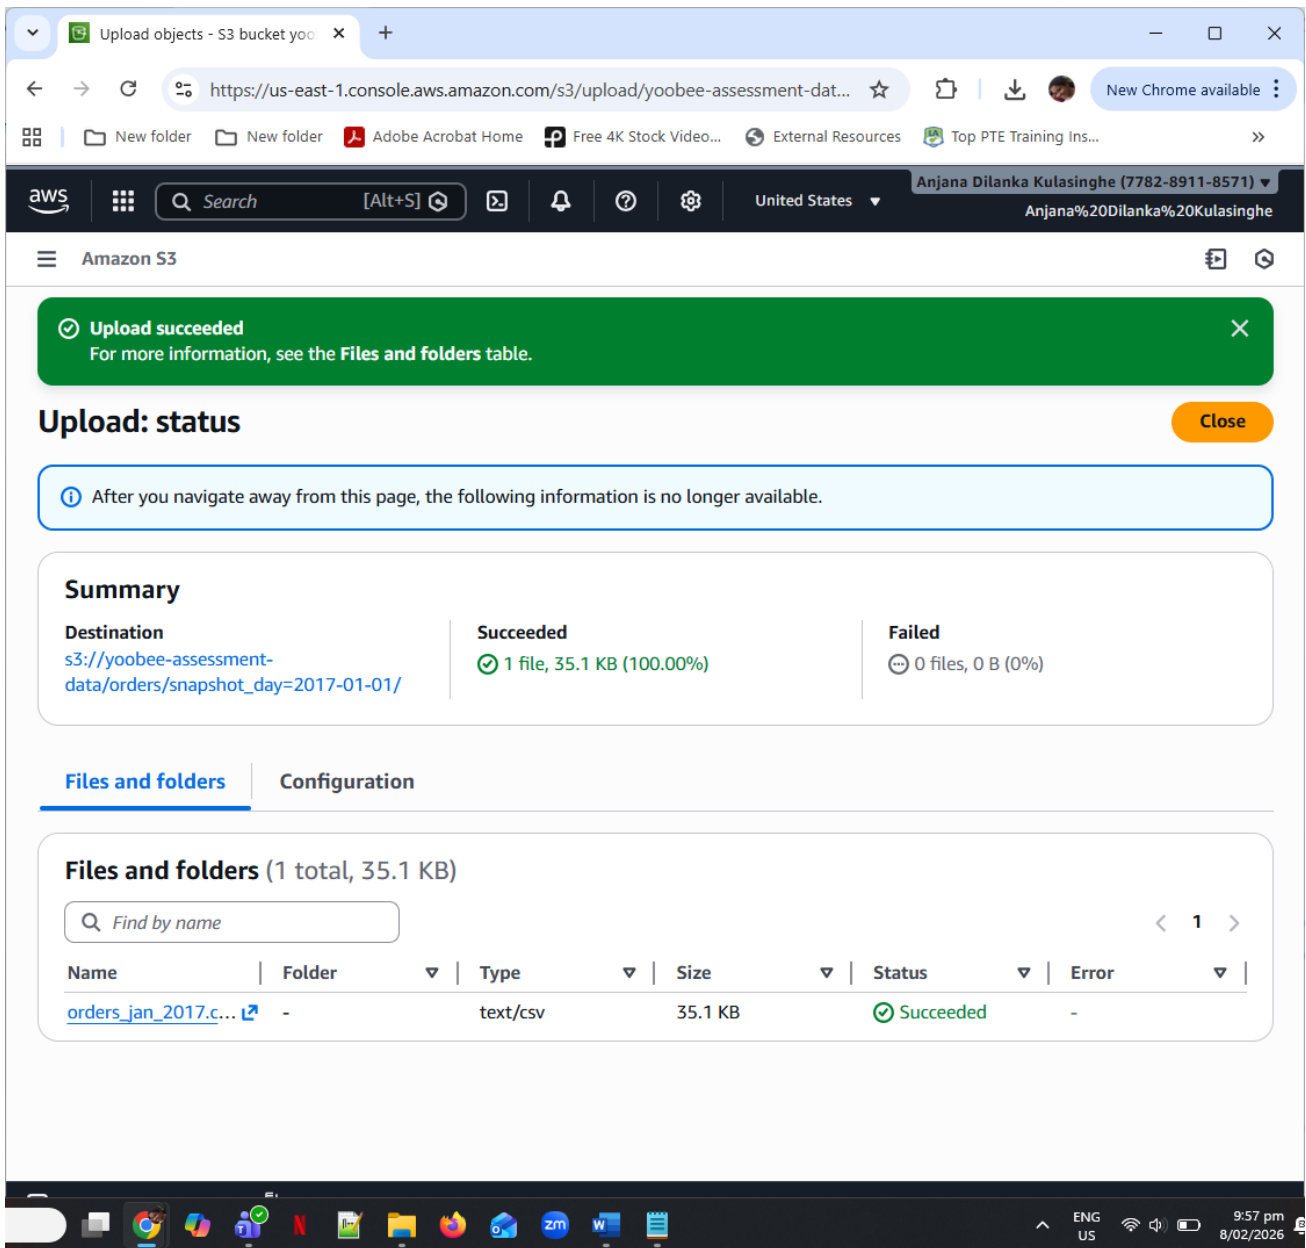
**Destination** Info

**Destination**  
[s3://yoobee-assessment-data/orders/snapshot\\_day=2017-01-01/](s3://yoobee-assessment-data/orders/snapshot_day=2017-01-01/)

► **Destination details**  
 Bucket settings that impact new objects stored in the specified destination.

► **Permissions**  
 Grant public access and access to other AWS accounts.

► **Properties**  
 Specify storage class, encryption settings, tags, and more.



## Activities: 1.2: AWS Glue Configuration

A Glue database named `db_yoozee` was created to manage metadata. A crawler named `yoozee_crawler` was configured to scan the `orders/` directory inside the S3 bucket. The crawler was assigned the appropriate IAM role to allow access to S3.

After creating an additional folder (`snapshot_day=2017-01-04`) and uploading another dataset, the crawler was re-run. After re-running the crawler, I could see the new partitions listed in the console. I did not need to manually edit the schema, which showed how Glue simplifies metadata management.

The screenshot shows the AWS Glue Databases console. The left sidebar contains navigation options for AWS Glue, including Getting started, ETL jobs, Data Catalog tables, and Data Catalog. The main content area displays 'Databases (1)' with a table listing the database 'db\_yoobee'. The table has columns for Name, Description, Location, Source catalog, and Created time. The database 'db\_yoobee' is listed with a description of '-' and a creation time of February 8, 2026.

Name	Description	Location	Source catalog	Created time
<a href="#">db_yoobee</a>	-	-	778289118571	February 8, 2026

**IAM Role "AWSGlueServiceRole-assessmentsql" successfully created**  
 Successfully created IAM Role "AWSGlueServiceRole-assessmentsql". This role trusts AWS Glue and has permissions to access your AWS Glue Crawler targets.

**Review and create**

**Step 1: Set crawler properties** [Edit](#)

**Set crawler properties**

Name	Description
yoobee_crawler	-

**Tags**  
-

**Step 2: Choose data sources and classifiers** [Edit](#)

**Data sources (1)** [Info](#)  
 The list of data sources to be scanned by the crawler.

Type	Data source	Parameters
S3	s3://yoobee-ass...	Recrawl all

**Step 3: Configure security settings** [Edit](#)

**Configure security settings**

IAM role	Security configuration
AWSGlueServiceRole-	-

The screenshot displays the AWS Glue console interface for a crawler named 'yoobee\_crawler'. At the top, a green notification banner states 'One crawler successfully created' with a success count of 2. The crawler's status is 'READY'. The 'Crawler properties' section lists the following details:

<b>Name</b> yoobee_crawler	<b>IAM role</b> AWSGlueServiceRole-assessmentsql	<b>Database</b> db_yoobee	<b>State</b> READY
<b>Description</b> -	<b>Security configuration</b> -	<b>Lake Formation configuration</b> -	<b>Table prefix</b> -
<b>Maximum table threshold</b> -			

Below the properties, there is an 'Advanced settings' section and a 'Crawler runs' section which currently shows 0 runs. The interface includes a left-hand navigation menu with categories like 'Data Catalog' and 'Data Integration and ETL', and a top navigation bar with the AWS logo and user information.

The screenshot displays the AWS Glue console interface for a crawler named 'yoobee\_crawler'. At the top, a green notification banner states 'Crawler successfully starting' with the message 'The following crawler is now starting: "yoobee\_crawler"'. Below this, the crawler's name 'yoobee\_crawler' is shown along with its last update time: 'February 8, 2026 at 09:06:48'. Action buttons for 'Run crawler', 'Edit', and 'Delete' are visible.

The 'Crawler properties' section is organized into a grid:

<b>Name</b> yoobee_crawler	<b>IAM role</b> AWSGlueServiceRole-assessmentsql	<b>Database</b> db_yoobee	<b>State</b> READY
<b>Description</b> -	<b>Security configuration</b> -	<b>Lake Formation configuration</b> -	<b>Table prefix</b> -
<b>Maximum table threshold</b> -			

Below the properties is an 'Advanced settings' section. At the bottom, the 'Crawler runs' section shows one run with options to 'Stop run', 'View CloudWatch logs', and 'View run details'. A search bar for 'Filter data' and a date/time filter are also present.

The screenshot displays the AWS Glue console interface for a crawler named 'yoobee\_crawler'. The left sidebar shows navigation options under 'AWS Glue' and 'Data Catalog'. The main content area shows the crawler's configuration, including its IAM role, database, and state (READY). Below the configuration, the 'Crawler runs' section shows a single completed run on February 8, 2026, with a duration of 01 min 02 s.

Name	IAM role	Database	State
yoobee_crawler	AWSGlueServiceRole-assessmentsql	db_yoobee	READY

Description	Security configuration	Lake Formation configuration	Table prefix
-	-	-	-

Maximum table threshold
-

**Crawler runs (1)**

The list of crawler runs for this crawler.

Start time (UTC)	End time (UTC)	Current/last duration	Status
February 8, 2026 at 09:...	February 8, 2026 at 09:...	01 min 02 s	Completed

The screenshot shows the AWS Glue console interface. At the top, there is a navigation bar with the AWS logo, a search bar, and the user's name 'Anjana Dilanka Kulasinghe'. The main content area is titled 'orders' and shows the table's details. A blue banner at the top of the console area announces new optimization features for Apache Iceberg tables. The table details section includes fields for Name, Classification, Database, Location, Description, and Last updated. The 'Table overview' tab is selected, and there are tabs for 'Schema', 'Partitions', 'Indexes', and 'Column statistics - new'. The footer of the console shows '© 2026, Amazon Web Services, Inc. or its affiliates.'

The screenshot shows the AWS Glue console interface. The breadcrumb navigation indicates the path: **AWS Glue** > **Tables** > **orders**. The main content area displays the **Schema (22)** for the 'orders' table, with the instruction 'View and manage the table schema.' Below this is a search bar labeled 'Filter schemas' and a table listing 16 columns.

#	Column name	Data type	Partition key	Comment
1	row id	bigint	-	-
2	order id	string	-	-
3	order date	string	-	-
4	ship date	string	-	-
5	ship mode	string	-	-
6	customer id	string	-	-
7	customer name	string	-	-
8	segment	string	-	-
9	country	string	-	-
10	city	string	-	-
11	state	string	-	-
12	postal code	bigint	-	-
13	region	string	-	-
14	product id	string	-	-
15	category	string	-	-
16	sub-category	string	-	-

The left sidebar contains navigation options under 'Data Catalog' (Databases, Tables, Stream schema registries, Schemas, Connections, Crawlers, Classifiers, Catalog settings) and 'Data Integration and ETL' (Zero-ETL integrations, ETL jobs, Visual ETL, Notebooks). The footer includes 'CloudShell', 'Feedback', 'Console Mobile App', and copyright information for Amazon Web Services, Inc. or its affiliates.

The screenshot displays the AWS Glue console interface. The browser address bar shows the URL: `https://us-east-1.console.aws.amazon.com/glue/home?region=us-east-1#/v2/data-catalog...`. The page title is "Table Detail - AWS Glue Console".

The left-hand navigation pane is expanded to "Data Catalog" > "Tables". The main content area shows the details for the table "orders".

**Table Details:**

- Name:** orders
- Classification:** CSV
- Database:** [db\\_yoobee](#)
- Location:** [s3://yoobee-assessment-data/orders/](#)
- Description:** -
- Connection:** -
- Last updated:** February 8, 2026 at 09:08:58

**Advanced properties:**

- Deprecated:** -
- Column statistics:** [No statistics](#)

**Partitions (1)**

The list of partitions for this table.

Filter partitions:

Partition Name	Files	Properties
2017-01-01	<a href="#">View files</a>	<a href="#">View Properties</a>

The footer of the console shows "© 2026, Amazon Web Services, Inc. or its affiliates." and links for "Privacy", "Terms", and "Cookie preferences".

The screenshot shows the Amazon S3 console interface. At the top, there is a navigation bar with the AWS logo, a search bar, and the user's name 'Anjana Dilanka Kulasinghe'. Below this, the main content area displays a green notification banner stating 'Upload succeeded' with a checkmark icon. Below the banner, there is a light blue information box that says 'After you navigate away from this page, the following information is no longer available.' The 'Summary' section shows the destination as 's3://yoobee-assessment-data/orders/snapshot\_day=2017-01-04/' and the upload status as 'Succeeded' with '1 file, 45.6 KB (100.00%)' and 'Failed' with '0 files, 0 B (0%)'. The 'Files and folders' section is active, showing a table with one file: 'orders\_apr\_2017.csv' (45.6 KB, text/csv, Succeeded). The bottom of the page includes a footer with 'CloudShell', 'Feedback', 'Console Mobile App', 'Privacy', 'Terms', 'Cookie preferences', and a copyright notice for Amazon Web Services, Inc. or its affiliates.

The screenshot displays the AWS Glue console interface. At the top, a green notification banner states: "Crawler successfully starting. The following crawler is now starting: 'yoobee\_crawler'". Below this, the "Crawlers" section is active, showing a table with one crawler entry:

Name	State	Schedule	Last run	Last run...	Log
yoobee_cra...	Running		February 8, 2026 at 09:19:14	Succeeded	View log

The left-hand navigation menu includes sections for "AWS Glue", "Data Catalog", and "Data Integration and ETL". The bottom of the page features a footer with "© 2026, Amazon Web Services, Inc. or its affiliates." and links for "Privacy", "Terms", and "Cookie preferences".

The screenshot shows the AWS Glue console interface. At the top, a green notification banner states: "Crawler successfully starting. The following crawler is now starting: 'yoobee\_crawler'". Below this, the "Crawlers" section is active, displaying a table with one crawler:

Name	State	Schedule	Last run	Last run...	Log
yoobee_cra...	Ready		Succeeded	February 8,...	<a href="#">View log</a>

The left sidebar contains navigation options for AWS Glue, including "Data Catalog" and "Data Integration and ETL". The bottom of the console shows "CloudShell", "Feedback", and "Console Mobile App" links.

The screenshot shows the "Table details" view for a table named "orders". The table is of type "CSV" and is located at "s3://yoobee-assessment-data/orders/". The "Partitions" tab is selected, showing two partitions:

Partition	Files	Properties
2017-01-01	<a href="#">View files</a>	<a href="#">View Properties</a>
2017-01-04	<a href="#">View files</a>	<a href="#">View Properties</a>

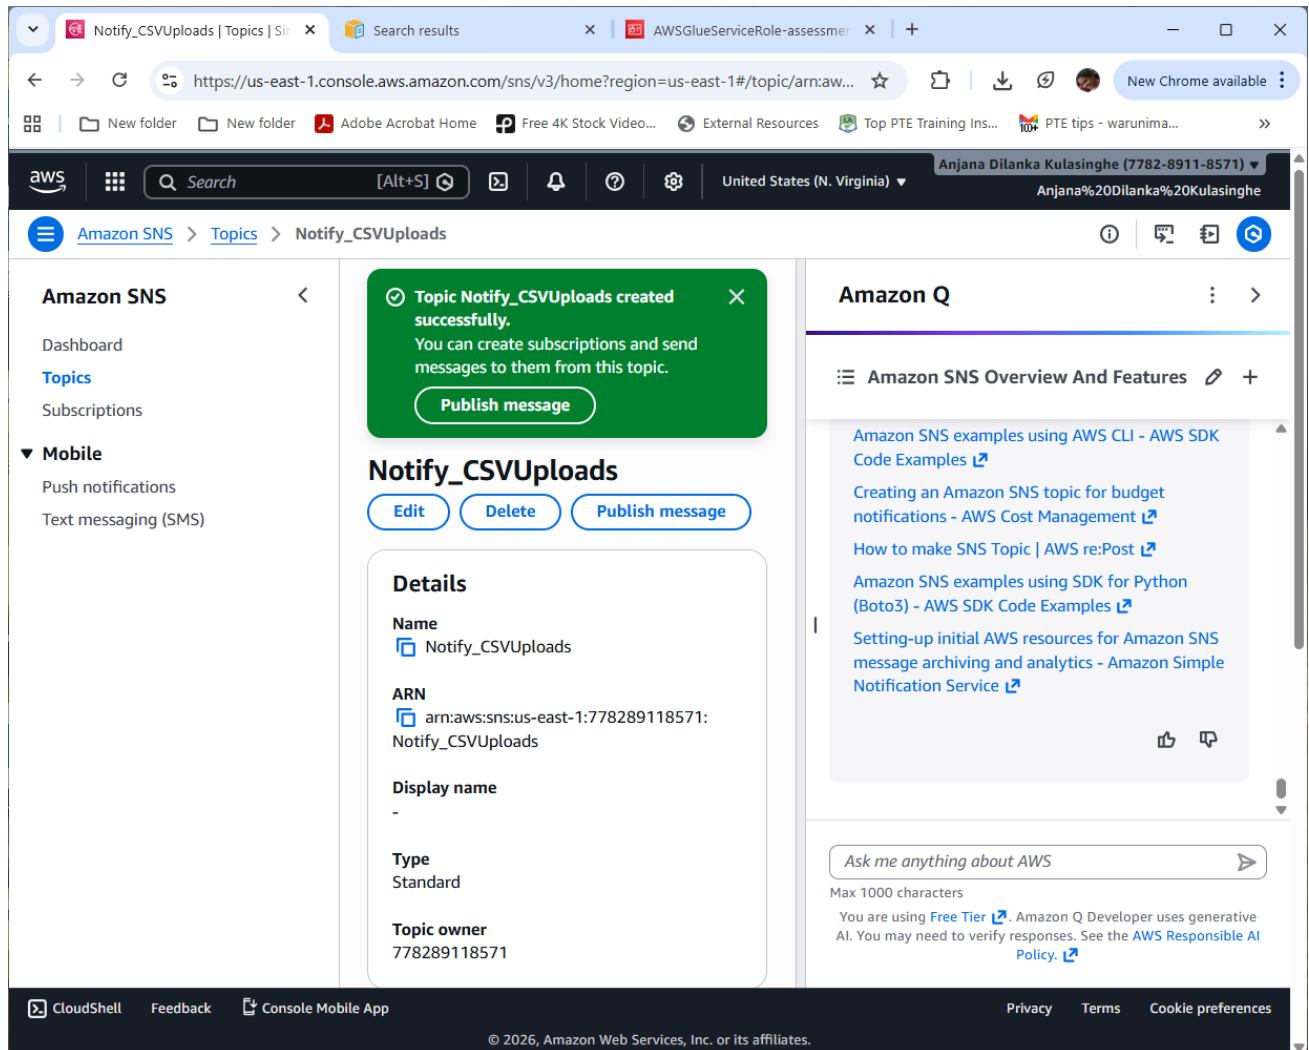
The left sidebar shows navigation options for "Data Integration and ETL" and "Legacy pages". The bottom of the console shows "CloudShell", "Feedback", and "Console Mobile App" links.

## Activities 1.3: Real-Time Notification Using AWS Lambda and SNS

To implement real-time monitoring, an SNS topic named Notify CSVUploads was created. A Lambda function was then deployed and configured with an S3 ObjectCreated trigger. The function checks whether the uploaded file has a .csv extension before publishing a structured notification message to the SNS topic.

This setup demonstrates event-driven automation. Instead of manually checking for new files, the system automatically reacts whenever new data is uploaded.

I tested this by uploading a new CSV file, and the notification was triggered immediately.



The screenshot displays the AWS Management Console interface for Amazon SNS. A green notification banner at the top states: "Subscription to Notify\_CSVUploads created successfully. The ARN of the subscription is arn:aws:sns:us-east-1:778289118571:Notify\_CSVUploads:282e2453-9943-47f5-a1d3-3306cf3c60a3." Below this, the subscription details are shown for "Subscription: 282e2453-9943-47f5-a1d3-3306cf3c60a3".

**Subscription: 282e2453-9943-47f5-a1d3-3306cf3c60a3**

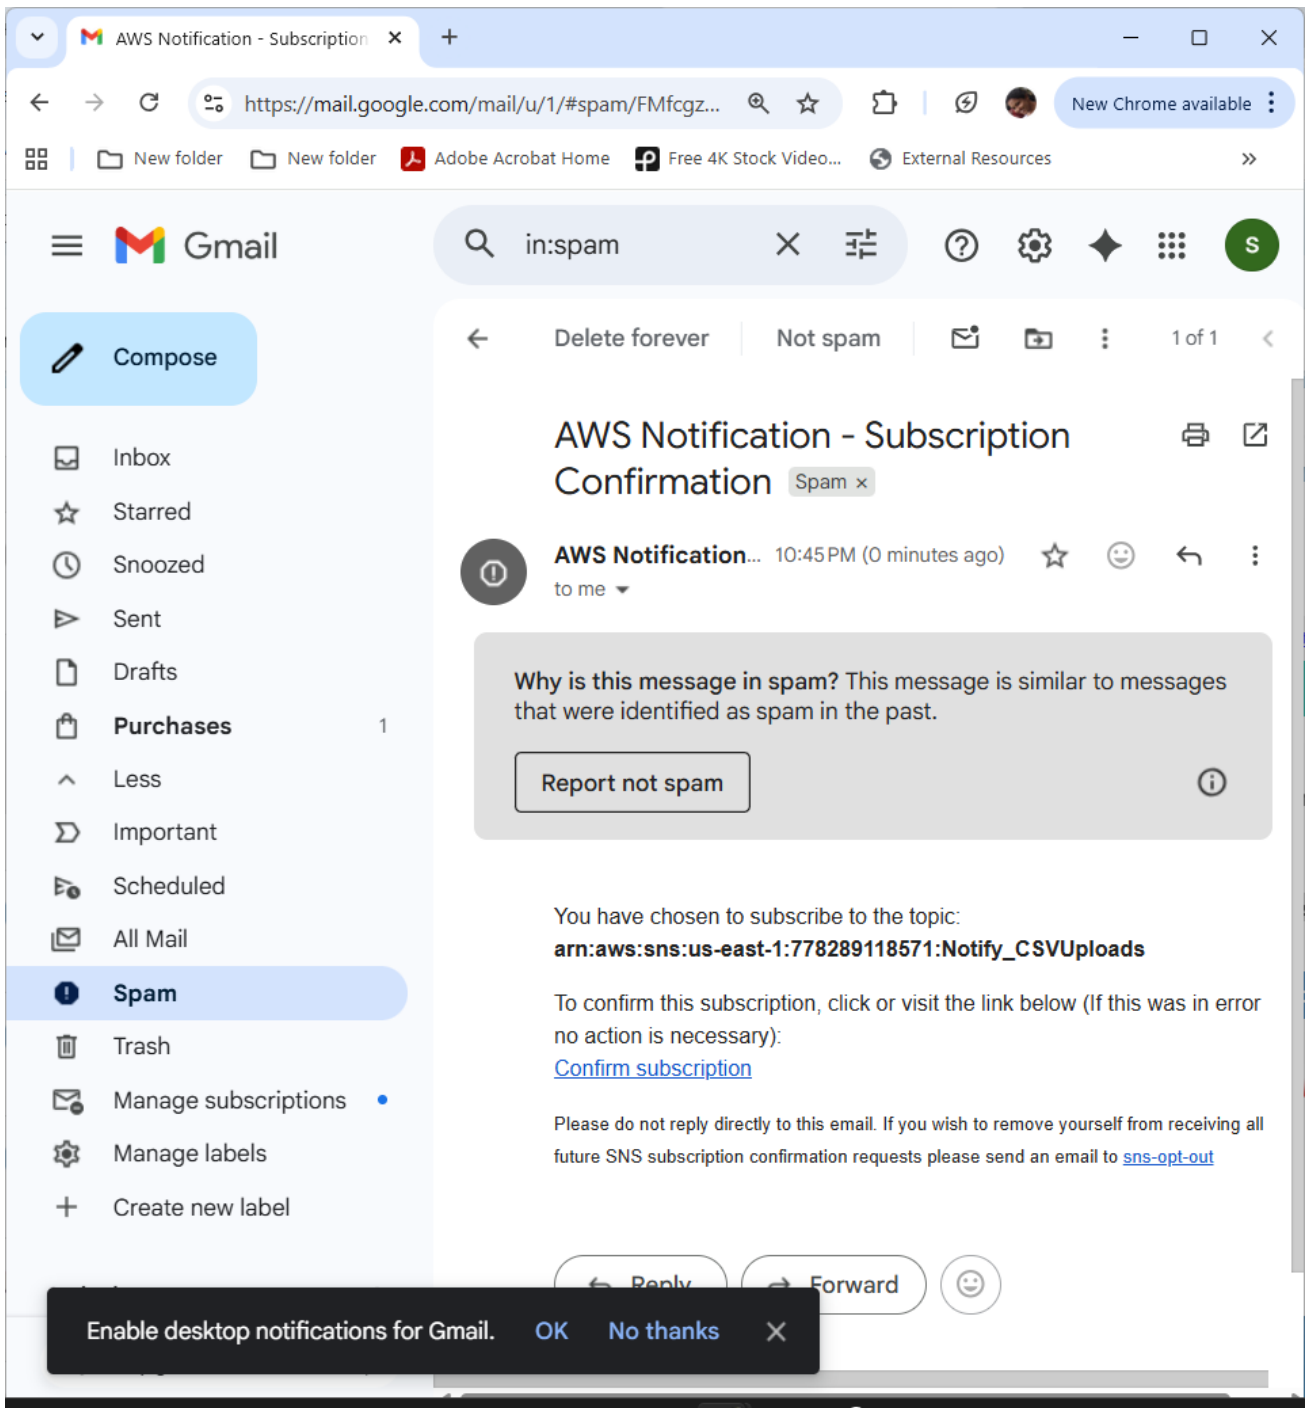
Buttons: [Edit](#) [Delete](#)

**Details**

<p><b>ARN</b></p> <p>arn:aws:sns:us-east-1:778289118571:Notify_CSVUploads:282e2453-9943-47f5-a1d3-3306cf3c60a3</p> <p><b>Endpoint</b></p> <p>genistudent97@gmail.com</p> <p><b>Topic</b></p> <p><a href="#">Notify_CSVUploads</a></p> <p><b>Subscription Principal</b></p> <p>arn:aws:iam::778289118571:root</p>	<p><b>Status</b></p> <p><span>🕒</span> Pending confirmation</p> <p><b>Protocol</b></p> <p>EMAIL</p>
--	---

Navigation: [Subscription filter policy](#) | [Redrive policy \(dead-letter queue\)](#)

Footer: CloudShell | Feedback | Console Mobile App | Privacy | Terms | Cookie preferences | © 2026, Amazon Web Services, Inc. or its affiliates.



The screenshot displays the AWS Lambda console interface for a function named 'csv\_upload\_notifier'. The 'Code source' tab is active, showing a Python file named 'lambda\_function.py'. The code is as follows:

```

1 import json
2 import boto3
3
4 sns = boto3.client('sns')
5
6 SNS_TOPIC_ARN = 'arn:aws:sns:us-east-1:778289118571:Notify_CSVUploads'
7
8 def lambda_handler(event, context):
9     print("Event received:", json.dumps(event))
10
11     record = event['Records'][0]
12     bucket_name = record['s3']['bucket']['name']
13     object_key = record['s3']['object']['key']
14
15     if object_key.endswith('.csv'):
16         message = f"CSV file uploaded: s3://{bucket_name}/{object_key}"
17
18         sns.publish(
19             TopicArn=SNS_TOPIC_ARN,
20             Subject="CSV Upload Notification",

```

The interface includes a left-hand sidebar with an Explorer view showing the file structure, a 'DEPLOY' section with 'Deploy (Ctrl+Shift+U)' and 'Test (Ctrl+Shift+I)' buttons, and a 'TEST EVENTS' section. The top navigation bar shows 'Lambda > Functions > csv\_upload\_notifier'. The bottom of the console features a footer with 'CloudShell', 'Feedback', 'Console Mobile App', 'Privacy', 'Terms', 'Cookie preferences', and a copyright notice for Amazon Web Services, Inc. or its affiliates.

The screenshot shows the AWS Lambda console interface. At the top, a green notification bar states "Successfully updated the function csv\_upload\_notifier." Below this, the "Code source" section is active, displaying the Python code for the lambda function. The code includes imports for 'json' and 'boto3', initializes an SNS client, and defines a lambda handler that triggers an SNS notification when a CSV file is uploaded to an S3 bucket.

**Code source Info** Open in Visual Studio Code Upload from

```

1  import json
2  import boto3
3
4  sns = boto3.client('sns')
5
6  SNS_TOPIC_ARN = 'arn:aws:sns:us-east-1:778289118571:Notify_CSVUploads'
7
8  def lambda_handler(event, context):
9      print("Event received:", json.dumps(event))
10
11     record = event['Records'][0]
12     bucket_name = record['s3']['bucket']['name']
13     object_key = record['s3']['object']['key']
14
15     if object_key.endswith('.csv'):
16         message = f"CSV file uploaded: s3://{bucket_name}/{object_key}"
17
18         sns.publish(
19             TopicArn=SNS_TOPIC_ARN,
20             Subject="CSV Upload Notification",

```

CloudShell Feedback Console Mobile App Privacy Terms Cookie preferences

© 2026, Amazon Web Services, Inc. or its affiliates.

The screenshot shows the AWS Lambda console for a function named 'csv\_upload\_notifier'. The interface is in the 'Configuration' tab. At the top, there are tabs for 'Diagram' and 'Template'. The 'Diagram' tab shows a visual representation of the function with an S3 bucket icon and a '+ Add trigger' button. To the right, a 'Description' panel shows the function's last modified time (16 minutes ago), its ARN (arn:aws:lambda:us-east-1:778289118571:function:csv\_upload\_notifier), and its URL. Below the diagram, there are tabs for 'Code', 'Test', 'Monitor', 'Configuration', 'Aliases', and 'Versions'. The 'Configuration' tab is active, showing a 'Triggers (1)' section with a search bar and a list of triggers. The first trigger is 'S3: yoobee-assessment-data' with the ARN 'arn:aws:s3:::yoobee-assessment-data'. A 'Details' link is provided for this trigger. The bottom of the console features a footer with 'CloudShell', 'Feedback', 'Console Mobile App', and copyright information for Amazon Web Services, Inc. (© 2026).

The screenshot shows the AWS S3 console interface. At the top, there is a navigation bar with the AWS logo, a search bar, and the region 'United States (N. Virginia)'. Below this, the page title is 'Amazon S3'. A green notification banner at the top left states 'Upload succeeded' with a close button. Below the notification, the 'Upload: status' section shows a message: 'After you navigate away from this page, the following information is no longer available.' The 'Summary' section displays the destination 's3://yoobee-assessment-data', 'Succeeded' status with '1 file, 35.1 KB (100.00%)', and 'Failed' status with '0 files, 0 B (0%)'. The 'Files and folders' tab is active, showing a table with one file: 'orders\_jan\_2017.csv' (text/csv, 35.1 KB, Succeeded). The footer contains links for CloudShell, Feedback, Console Mobile App, and copyright information for Amazon Web Services, Inc. (2026).

The screenshot shows the AWS CloudWatch console interface. The browser address bar indicates the URL: `https://us-east-1.console.aws.amazon.com/cloudwatch/home?region=us-east-1#logsV2:log-groups/log-grou...`. The console header shows the user is logged in as Anjana Dilanka Kulasinghe (7782-8911-8571) in the United States (N. Virginia) region.

The main content area is titled `/aws/lambda/csv_upload_notifier`. It features a navigation sidebar on the left with categories like Alarms, AI Operations, GenAI Observability, Application Signals (APM), Infrastructure Monitoring, and Logs. The `Log Management` option is selected.

The central panel displays the **Log group details** for the group `/aws/lambda/csv_upload_notifier`. The details are organized into several sections:

- Log class:** Standard
- ARN:** `arn:aws:logs:us-east-1:778289118571:log-group:/aws/lambda/csv_upload_notifier:*`
- Creation time:** -
- Retention:** Never expire
- Stored bytes:** -
- Account:** -
- Metric filters:** Not supported
- Subscription filters:** Not supported
- Contributor Insights rules:** Not supported
- KMS key ID:** -
- Deletion protection:** Off
- Data protection:** -
- Sensitive data count:** Not supported
- Custom field indexes:** Not supported
- Transformer:** Not supported
- Anomaly detection:** Not supported

At the bottom of the console, there are navigation tabs for `Log streams`, `Tags`, `Data protection`, `Anomaly detection`, `Metric filters`, and `Subscription filters`. The footer contains the text: © 2026, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences.

## Task 2: Query and Visualize the Dataset Using Athena and QuickSight.

In this task, Amazon Athena was used to query the dataset stored in S3, and Amazon QuickSight was used to visualize analytical insights. This stage transformed raw data into meaningful business intelligence.

### Activities: 2.1: Amazon Athena Queries

An S3 bucket was configured to store Athena query results. Athena was connected to the Glue Data Catalog, allowing SQL queries to be executed directly on the orders table.

Queries were executed to calculate total sales by category, retrieve full dataset records, and filter by specific snapshot dates. When a second dataset was uploaded and the crawler re-run, the data scanned size increased, which indicated that the additional dataset had been successfully included in the table.

**Successfully created bucket "athena-queryresults-bucket1"**  
To upload files and folders, or to configure additional bucket settings, choose [View details](#).

**General purpose buckets** All AWS Regions | **Directory buckets**

**General purpose buckets (2)** [Info](#) [Refresh](#) [Copy ARN](#) [Empty](#) [Delete](#) [Create bucket](#)

Buckets are containers for data stored in S3.

Name	AWS Region	Creation date
<a href="#">athena-queryresults-bucket1</a>	US East (N. Virginia) us-east-1	February 10, 2026, 11:43:49 (UTC+13:00)
<a href="#">yoobee-assessment-data</a>	US East (N. Virginia) us-east-1	February 8, 2026, 21:52:08 (UTC+13:00)

**Account snapshot** [Info](#) [View dashboard](#)  
Updated daily  
Storage Lens provides visibility into storage usage and activity trends.

**External access summary** [Info](#)  
Updated daily  
External access findings help you identify bucket permissions that allow public access or access from other AWS accounts.

CloudShell Feedback Console Mobile App Privacy Terms Cookie preferences

© 2026, Amazon Web Services, Inc. or its affiliates.

Query editor | Athena | us-east-1

https://us-east-1.console.aws.amazon.com/athena/home?region=us-east-1#/query-editor/settings

Settings successfully updated.

**Query in Amazon SageMaker Unified Studio** Set up

Use your IAM role to analyze and build with your existing AWS resources in a single data and AI development environment.

Editor | Recent queries | Saved queries | **Query settings** | Workgroup: primary

### Query settings

Query settings temporarily override your workgroup configurations. Clear all the query settings to enforce your workgroup configurations.

#### Query result encryption

**Query result location**  
s3://athena-queryresults-bucket1/

**Encrypt query results**  
-

**Expected bucket owner**  
-

**Assign bucket owner full control over query results**  
Turned off

[Manage](#)

#### Execution controls Info

Set limits and control query execution behavior for the workgroup. Controls apply to all queries in the workgroup unless overridden.

Search execution controls

Control Name	Value	Unit
Data scanned limit	-	-

[Edit controls](#)

CloudShell | Feedback | Console Mobile App | Privacy | Terms | Cookie preferences

© 2026, Amazon Web Services, Inc. or its affiliates.

The screenshot displays the Amazon Athena Query Editor interface. On the left, the 'Data' sidebar is visible, showing the following configuration:

- Data source:** AwsDataCatalog
- Catalog:** None
- Database:** db\_yoobee
- Tables and views:** A search filter is present, and a list of tables is shown under 'Tables (1)'. The 'orders' table is selected, with columns: row id (bigint), order id (string), order date (string), ship date (string), ship mode (string), customer id (string), customer name (string), and segment (string).

The main workspace shows 'Query 1' with a SQL editor containing 'Ln 1, Col 1'. Below the editor are buttons for 'Run', 'Explain', 'Cancel', 'Clear', and 'Create'. A 'Reuse query results' toggle is also present. The 'Query results' section shows 'No results' and a message 'Run a query to view results'.

The screenshot displays the Amazon Athena Query Editor interface. On the left, the 'Data' sidebar shows configuration for the 'AwsDataCatalog', 'None' catalog, and 'db\_yoobee' database. Below this, a table named 'orders' is listed with columns: row id (bigint), order id (string), order date (string), ship date (string), ship mode (string), customer id (string), customer name (string), and segment (string). The main editor area contains a SQL query:

```

1 SELECT category, SUM(sales)
2 FROM "db_yoobee"."orders"
3 GROUP BY category;
4

```

Below the query editor, the 'Query results' tab is active, showing a 'Completed' status with the following metrics: Time in queue: 102 ms, Run time: 766 ms, and Data scanned: 80.68 KB. The results table has 3 rows and 2 columns: # and category.

The screenshot displays the Amazon Athena Query Editor interface. On the left, a sidebar shows the database 'db\_yoobee' and a table named 'orders' with various columns and their data types. The main area shows the SQL editor with a query executed. Below the editor, the 'Query results' tab is active, displaying a completion status and a table of results.

**Query Results:**

#	category	_col1
1	Technology	28721.978000000003
2	Office Supplies	34334.544999999999
3	Furniture	7627.393099999999

The screenshot displays the AWS Athena Query Editor interface. On the left, the 'Data' sidebar shows configuration for the 'AwsDataCatalog' data source, 'None' catalog, and 'db\_yoobee' database. Below this, a list of tables is shown, with 'orders' selected. The 'orders' table schema includes columns: row id (bigint), order id (string), order date (string), ship date (string), ship mode (string), customer id (string), customer name (string), and segment (string).

The main editor area shows a SQL query:
 

```

    1 SELECT *
    2 FROM "db_yoobee"."orders";
    3
    
```

 Below the query editor, there are buttons for 'Run again', 'Explain', 'Cancel', 'Clear', and 'Create'. A 'Reuse query results' toggle is also present.

The 'Query results' section shows a 'Completed' status with the following metrics:
 

- Time in queue: 103 ms
- Run time: 587 ms
- Data scanned: 80.68 KB

 There are buttons for 'Copy' and 'Download results CSV'. Below this, a search bar for 'Search rows' is visible, and the top of a table with columns: #, row id, order id, order date, ship date, ship mode, and customer id is shown.

The screenshot displays the Amazon Athena Query Editor interface. On the left, a table list shows the 'orders' table with columns: row id (bigint), order id (string), order date (string), ship date (string), ship mode (string), customer id (string), customer name (string), segment (string), country (string), city (string), state (string), and postal code (bigint). The main area shows a query execution result for the 'orders' table. The query is in SQL, and the results are displayed in a table with 7 rows and 6 columns: #, row id, order id, order date, ship date, ship mode, and customer id. The query status is 'Completed' with a time in queue of 103 ms, a run time of 587 ms, and 80.68 KB of data scanned. The results table contains the following data:

#	row id	order id	order date	ship date	ship mode	customer id
1	13	CA-2017-114412	2017-04-15	4/20/2017	Standard Class	AA-10480
2	177	US-2017-152366	2017-04-21	4/25/2017	Second Class	SJ-20500
3	232	US-2017-100930	2017-04-07	4/12/2017	Standard Class	CS-12400
4	233	US-2017-100930	2017-04-07	4/12/2017	Standard Class	CS-12400
5	234	US-2017-100930	2017-04-07	4/12/2017	Standard Class	CS-12400
6	235	US-2017-100930	2017-04-07	4/12/2017	Standard Class	CS-12400
7	236	US-2017-100930	2017-04-07	4/12/2017	Standard Class	CS-12400

The screenshot displays the Amazon Athena Query Editor interface. At the top, the browser address bar shows the URL: `https://us-east-1.console.aws.amazon.com/athena/home?region=us-east-1#/query-editor/history/ec7cf854-c3e8-4a76-9503-9cf03b9a...`. The AWS navigation bar includes the user name 'Anjana Dilanka Kulasinghe (7782-8911-8571)' and the region 'United States (N. Virginia)'. The main header shows 'Amazon Athena > Query editor' and a 'Workgroup' dropdown set to 'primary'. A notification banner states: 'Athena now supports typeahead code suggestions to speed up SQL query development. Typeahead suggestions are turned on by default. You can change this setting in query editor preferences.' The left sidebar, titled 'Data', contains configuration options for 'Data source' (set to 'AwsDataCatalog'), 'Catalog' (set to 'None'), and 'Database' (set to 'db\_yoobee'). Below these are 'Tables and views' with a search filter and a list of tables. The 'Tables (1)' section shows a table named 'orders' with columns 'row id' (bigint) and 'order id' (string). The main editor area shows three query tabs, with 'Query 3' selected. The SQL query is:
 

```

1 SELECT *
2 FROM "db_yoobee"."orders"
3 WHERE snapshot_day = '2017-01-01';
4
    
```

 Below the query editor, the status 'SQL Ln 2, Col 16' is shown. Action buttons include 'Run again', 'Explain', 'Cancel', 'Clear', and 'Create'. A 'Reuse query results' toggle is also present, indicating results are available 'up to 60 minutes ago'. The footer contains 'CloudShell', 'Feedback', 'Console Mobile App', and copyright information for Amazon Web Services, Inc.

The screenshot displays the Amazon Athena Query Editor interface. On the left, a schema view lists columns such as row id, order id, order date, ship date, ship mode, customer id, customer name, segment, country, city, state, and postal code. The main area shows a query execution summary: 'Completed', 'Time in queue: 106 ms', 'Run time: 553 ms', and 'Data scanned: 35.11 KB'. Below this, a table of 155 results is displayed, with the first 10 rows visible. The table has columns for row id, order id, order date, ship date, ship mode, and customer id.

#	row id	order id	order date	ship date	ship mode	customer id
1	440	CA-2017-157252	2017-01-20	1/23/2017	Second Class	CV-12805
2	516	CA-2017-127432	2017-01-22	1/27/2017	Standard Class	AD-10180
3	517	CA-2017-127432	2017-01-22	1/27/2017	Standard Class	AD-10180
4	518	CA-2017-127432	2017-01-22	1/27/2017	Standard Class	AD-10180
5	519	CA-2017-127432	2017-01-22	1/27/2017	Standard Class	AD-10180
6	523	CA-2017-145142	2017-01-23	1/25/2017	First Class	MC-17605
7	733	CA-2017-131954	2017-01-21	1/25/2017	Standard Class	DS-13030
8	734	CA-2017-131954	2017-01-21	1/25/2017	Standard Class	DS-13030
9	735	CA-2017-131954	2017-01-21	1/25/2017	Standard Class	DS-13030
10	736	CA-2017-131954	2017-01-21	1/25/2017	Standard Class	DS-13030

## Activities: 2.2: Amazon QuickSight Dashboards

Amazon QuickSight was configured using the Standard Edition and connected to Athena. SPICE was used to import the dataset for faster visual performance.

From the chart, it was clear that certain categories generated significantly higher sales, which would help decision-makers prioritise products.

A bar chart was created to show Sales by Category, and another visual was built to represent Profit by City. These dashboards transformed numerical results into visual insights that would support business decision-making.

**QUICK SUITE ACCESS TO AWS SERVICES**

Make your existing AWS data and users available in Quick Suite. [Learn more](#)

**IAM Role**

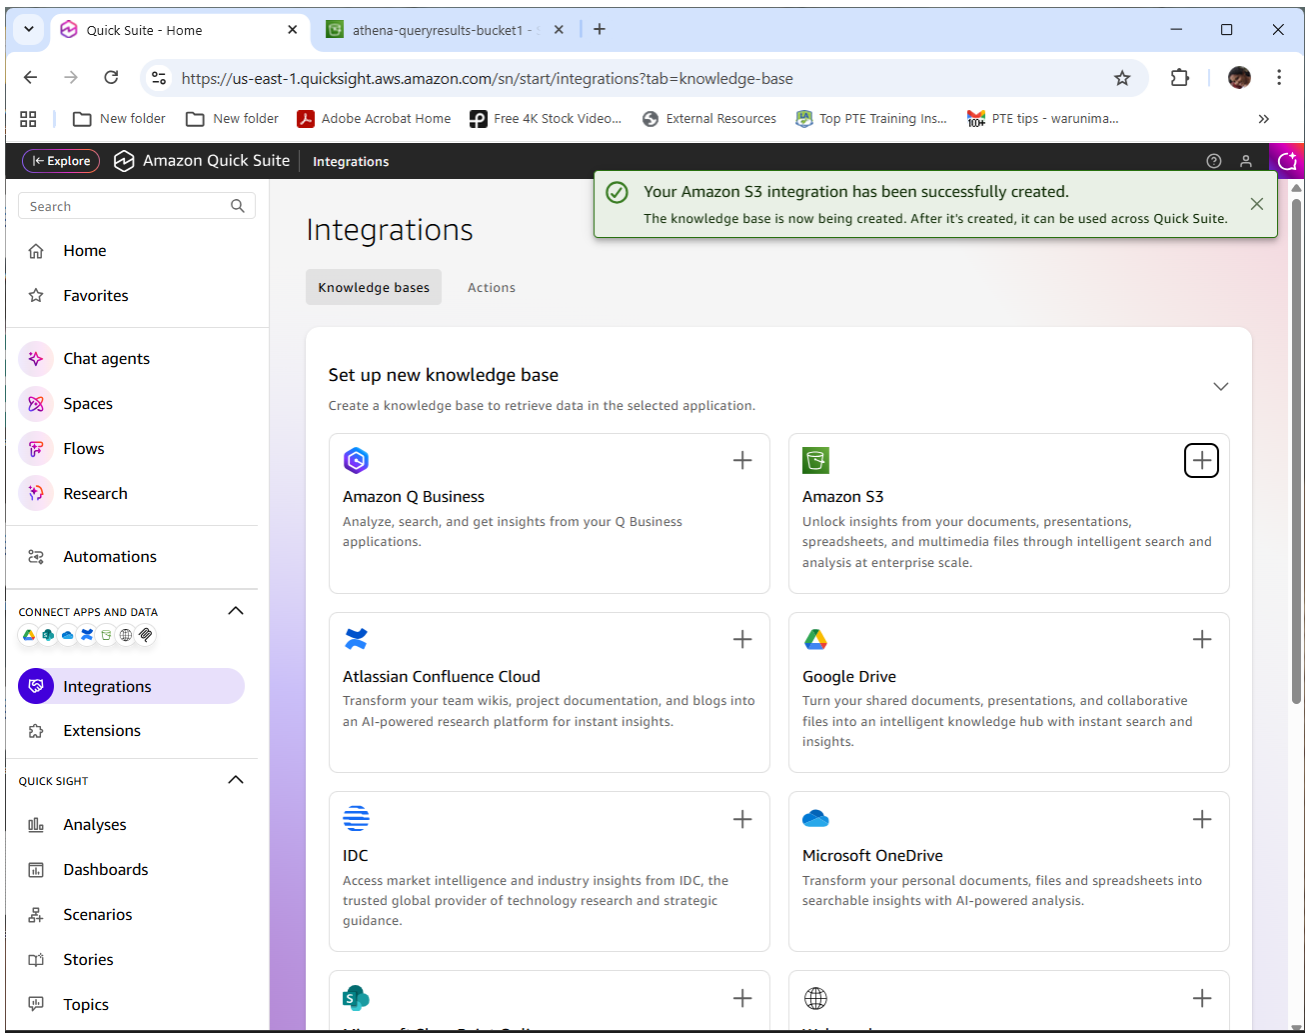
- Use Quick Suite-managed role (default)
- Use an existing role

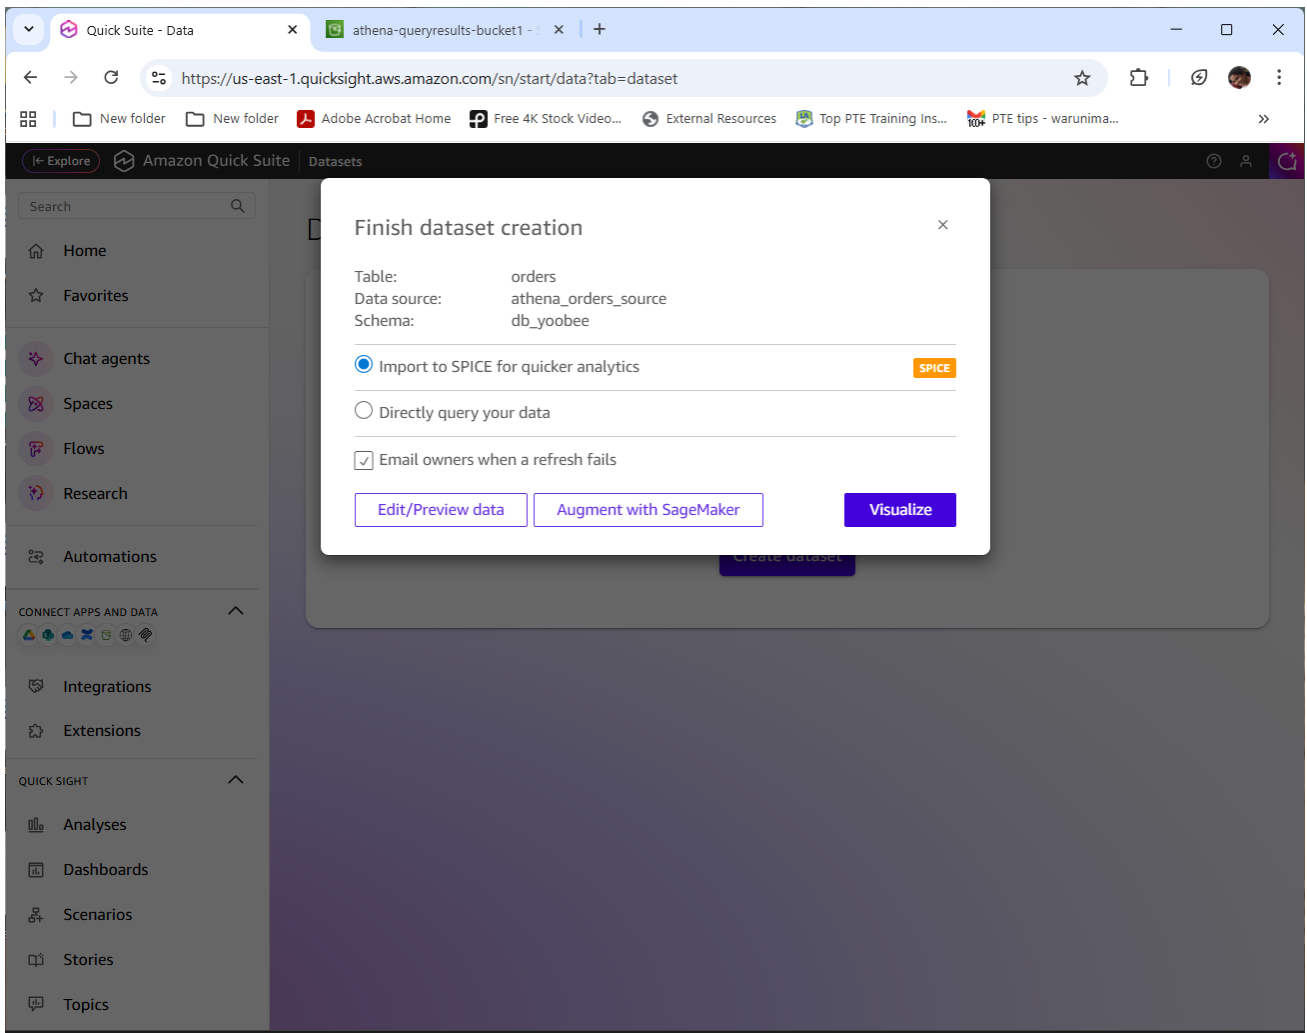
**Create or access Amazon Q Business applications**

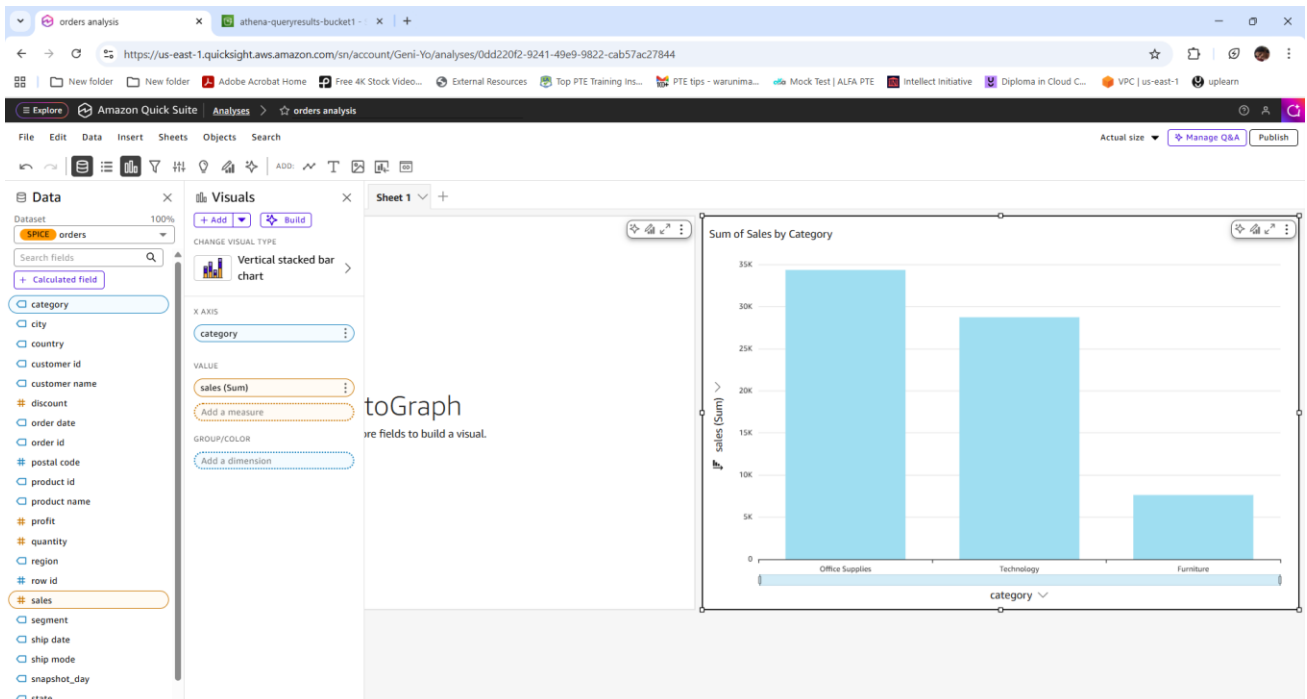
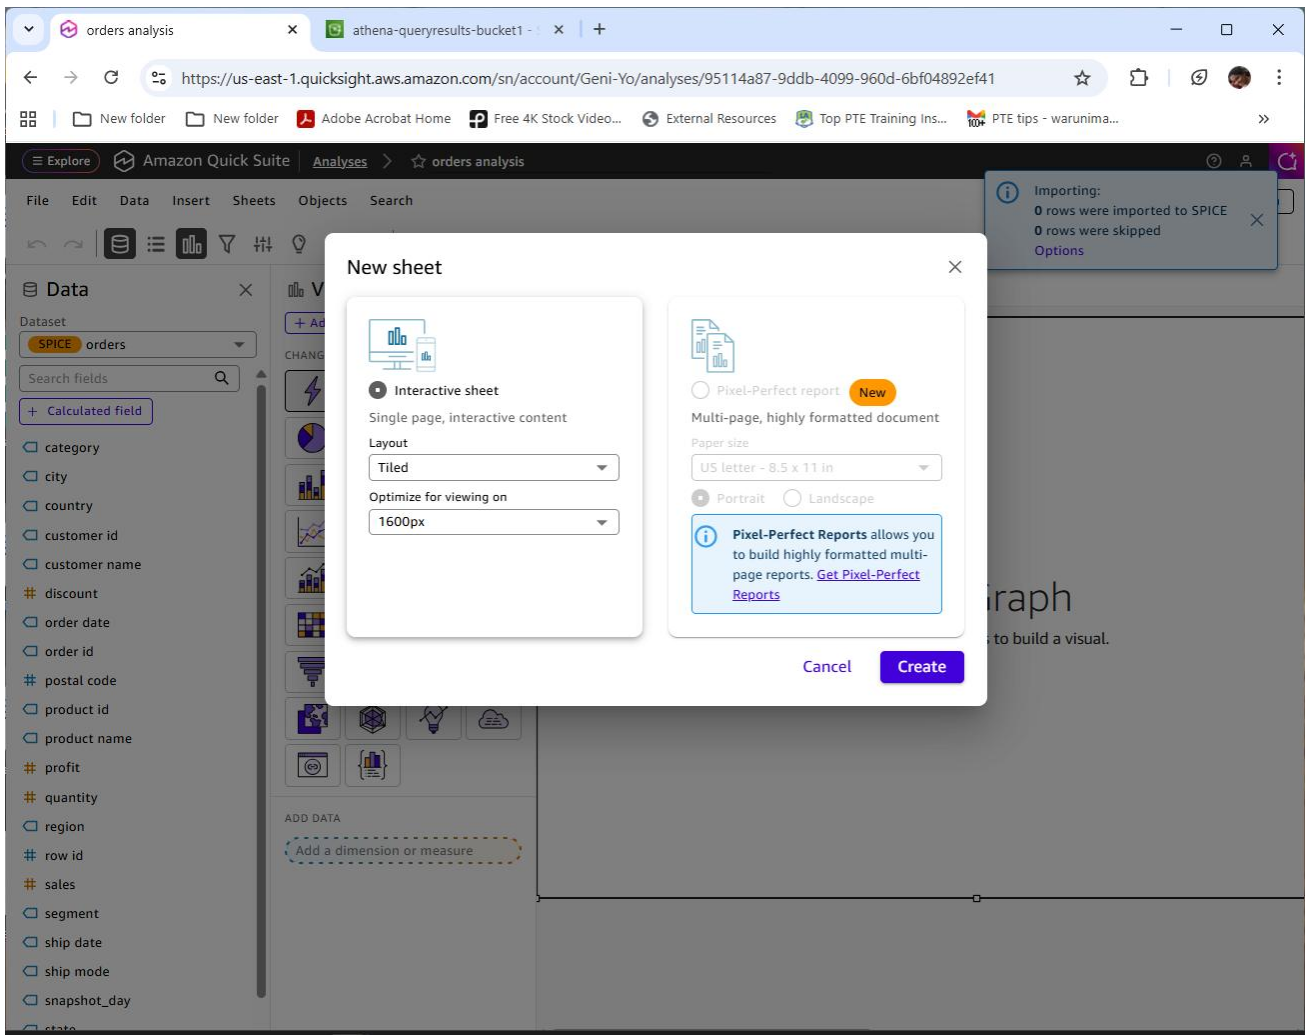
- Amazon Q Business  
Link Amazon Q Business applications to find answers from your existing Q Business indices.  
[Select application](#)

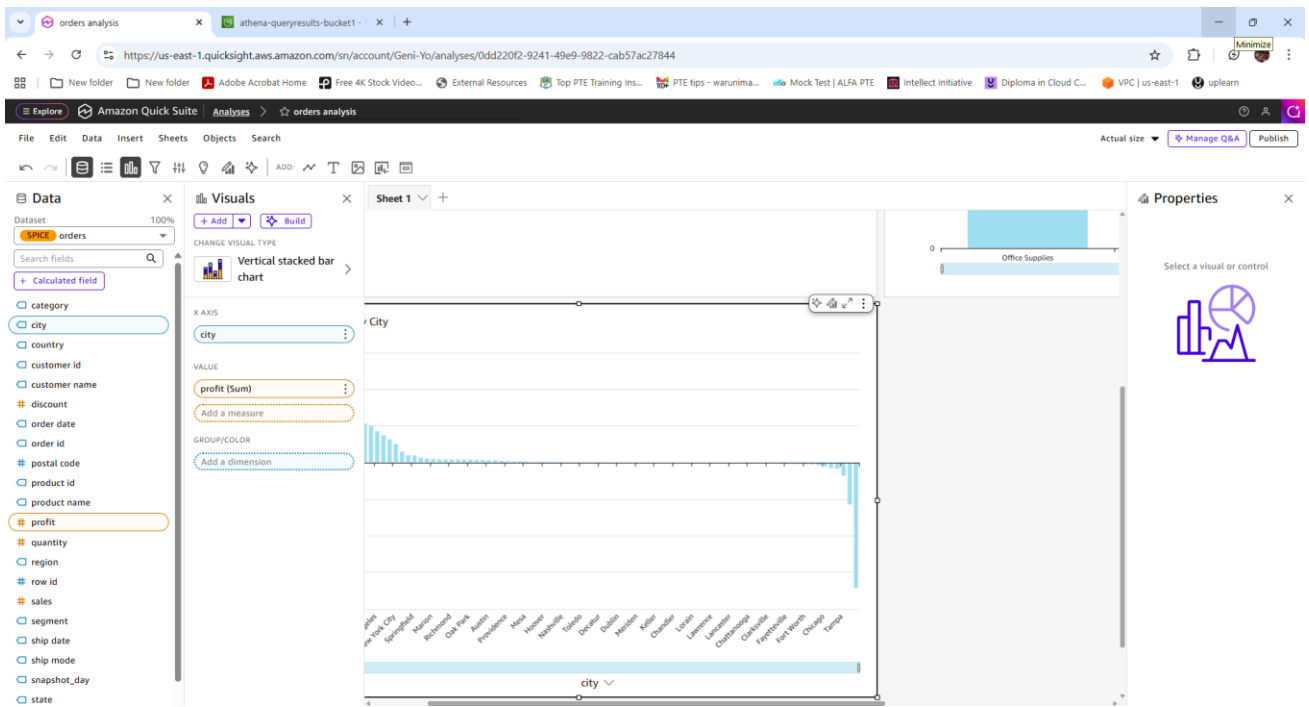
**Allow access and autodiscovery for these resources**

- Amazon Redshift
- Amazon RDS
- IAM
- Amazon S3 (2 buckets selected)  
[Select S3 buckets](#)
- Amazon Athena  
Make sure you've chosen the right Amazon S3 buckets for Quick Suite access
- Amazon S3 Storage Analytics
- Amazon OpenSearch Service
- Amazon SageMaker
- Amazon Timestream
- AWS SecretsManager  
[Select secrets](#)
- AWS SecretsManager Write







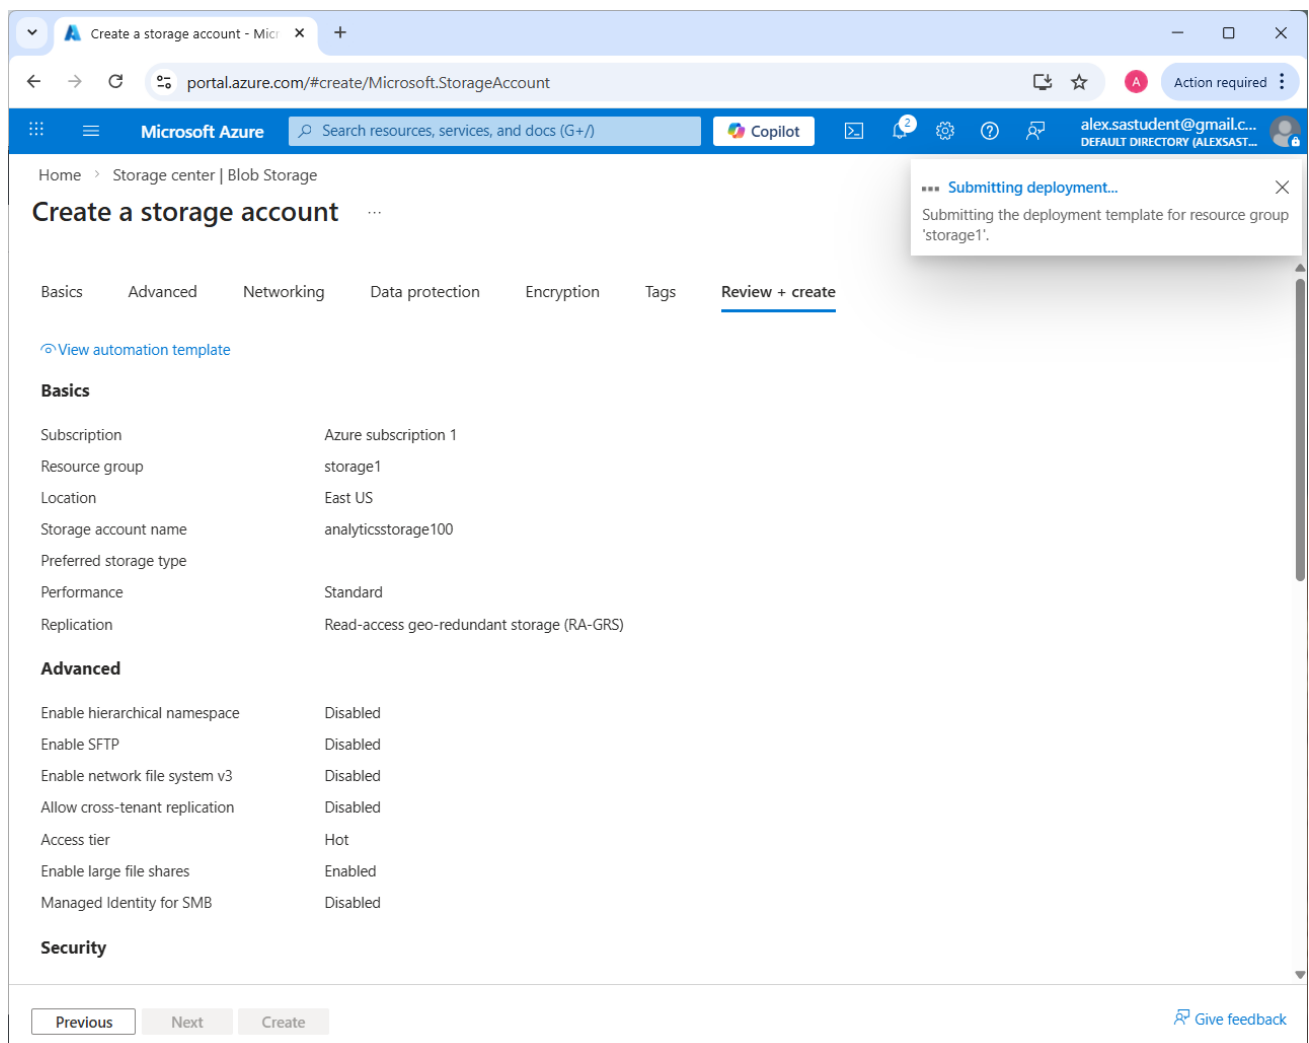


## Task 3: Multi-Cloud Integration (GitHub → AWS → Azure Synapse Analytics)

This task demonstrates cross-cloud interoperability by exporting curated AWS query results and analysing them within Microsoft Azure Synapse Analytics. This reflects real-world enterprise scenarios where organisations operate across multiple cloud providers.

### Activity 3.1: Data Export from AWS S3

The Athena query results were downloaded from the S3 result location and renamed for clarity. These CSV files represented curated datasets ready for cross-cloud integration.



Resource Manager - Microsoft

portal.azure.com/?ocid=AIDcmm3m06wb2\_SEM\_k\_f6ffa487948117ca6fdec0ba9a76752\_k\_#view/HubsExten...

Microsoft Azure Search resources, services, and docs (G+)

Home > Resource Manager

### Resource Manager | Resource groups

Default Directory

Search

+ Create Manage view Refresh Export to CSV Open query Assign tags Group by none

Filter for any field... Subscription equals all Location equals all Add filter

<input type="checkbox"/>	Name ↑	Subscription	Location
<input type="checkbox"/>	rg-lab	Azure subscription 1	Austria East

Showing 1 - 1 of 1. Display count: auto

Add or remove favorites by pressing Ctrl+Shift+F

Give feedback

The screenshot shows the Microsoft Azure portal interface. At the top, the browser address bar displays the URL: `portal.azure.com/?ocid=AIDcmmp3m06wb2_SEM_k_f6ffa487948117ca6fdec0ba9a76752_k_#view/HubsExten...`. The page title is "labstorage010\_1770799327171 | Overview".

The main content area features a navigation pane on the left with options: Overview (selected), Inputs, Outputs, and Template. The central pane displays a green checkmark icon and the message: "Your deployment is complete". Below this, deployment details are listed:
 

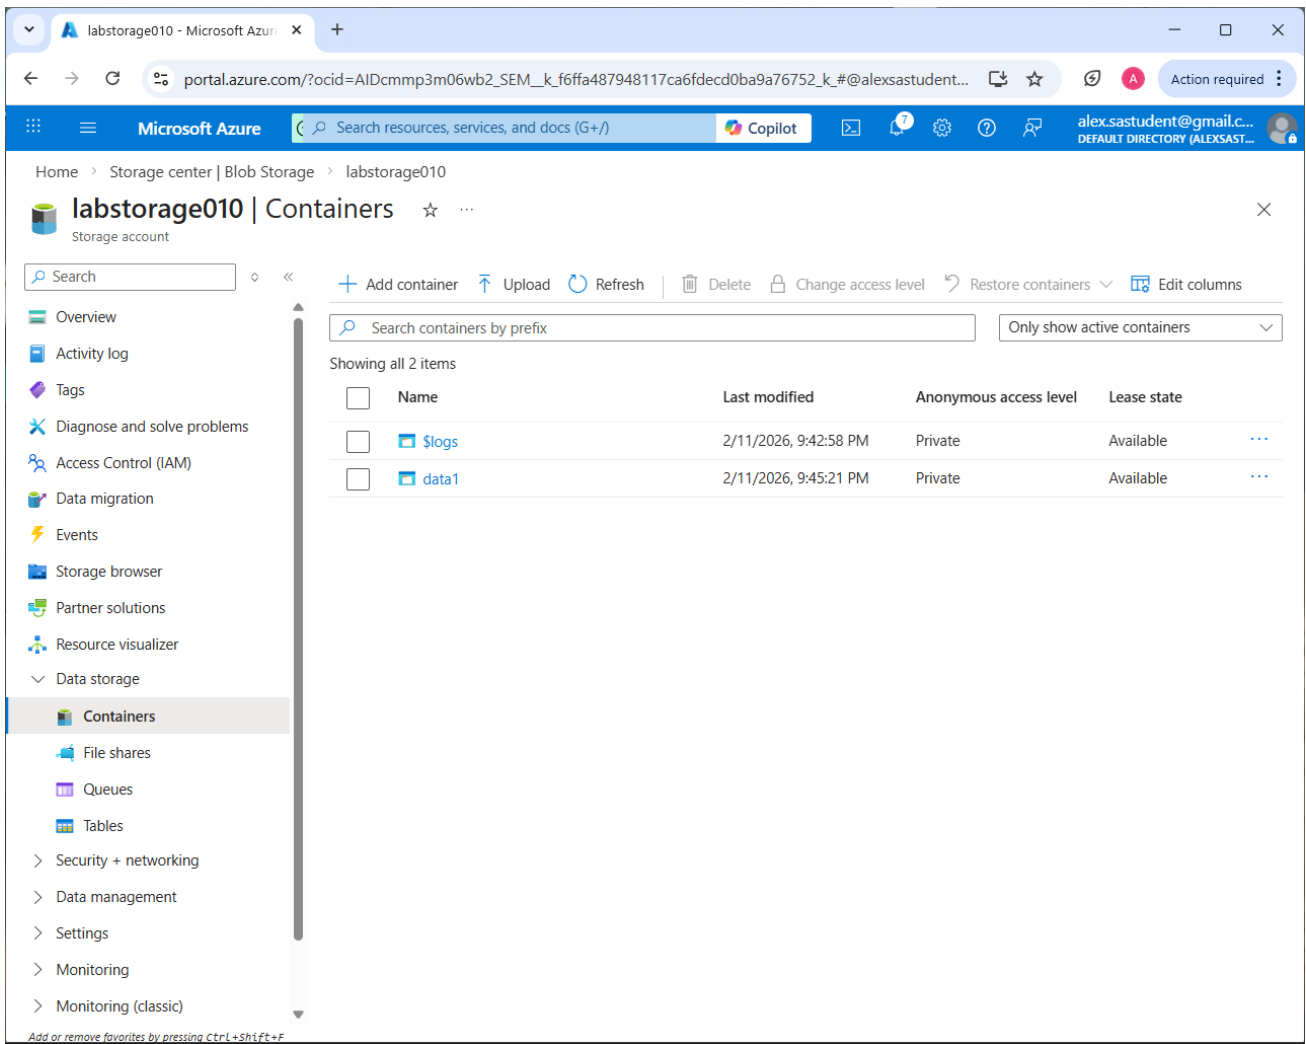
- Deployment name: labstorage01...
- Subscription: Azure subscription 1
- Resource group: rg-lab
- Start time: 2/11/2026, 9:42:21 PM
- Correlation ID: 19959987-8b5a-4004-9460-0

 The right-hand side of the page contains several recommendations:
 

- Cost Management:** Get notified to stay within your budget and prevent unexpected charges on your bill. [Set up cost alerts >](#)
- Microsoft Defender for Cloud:** Secure your apps and infrastructure. [Go to Microsoft Defender for Cloud >](#)
- Free Microsoft tutorials:** [Start learning today >](#)
- Work with an expert:** Azure experts are service provider partners who can help manage your assets on Azure and be your first line of support. [Find an Azure expert >](#)

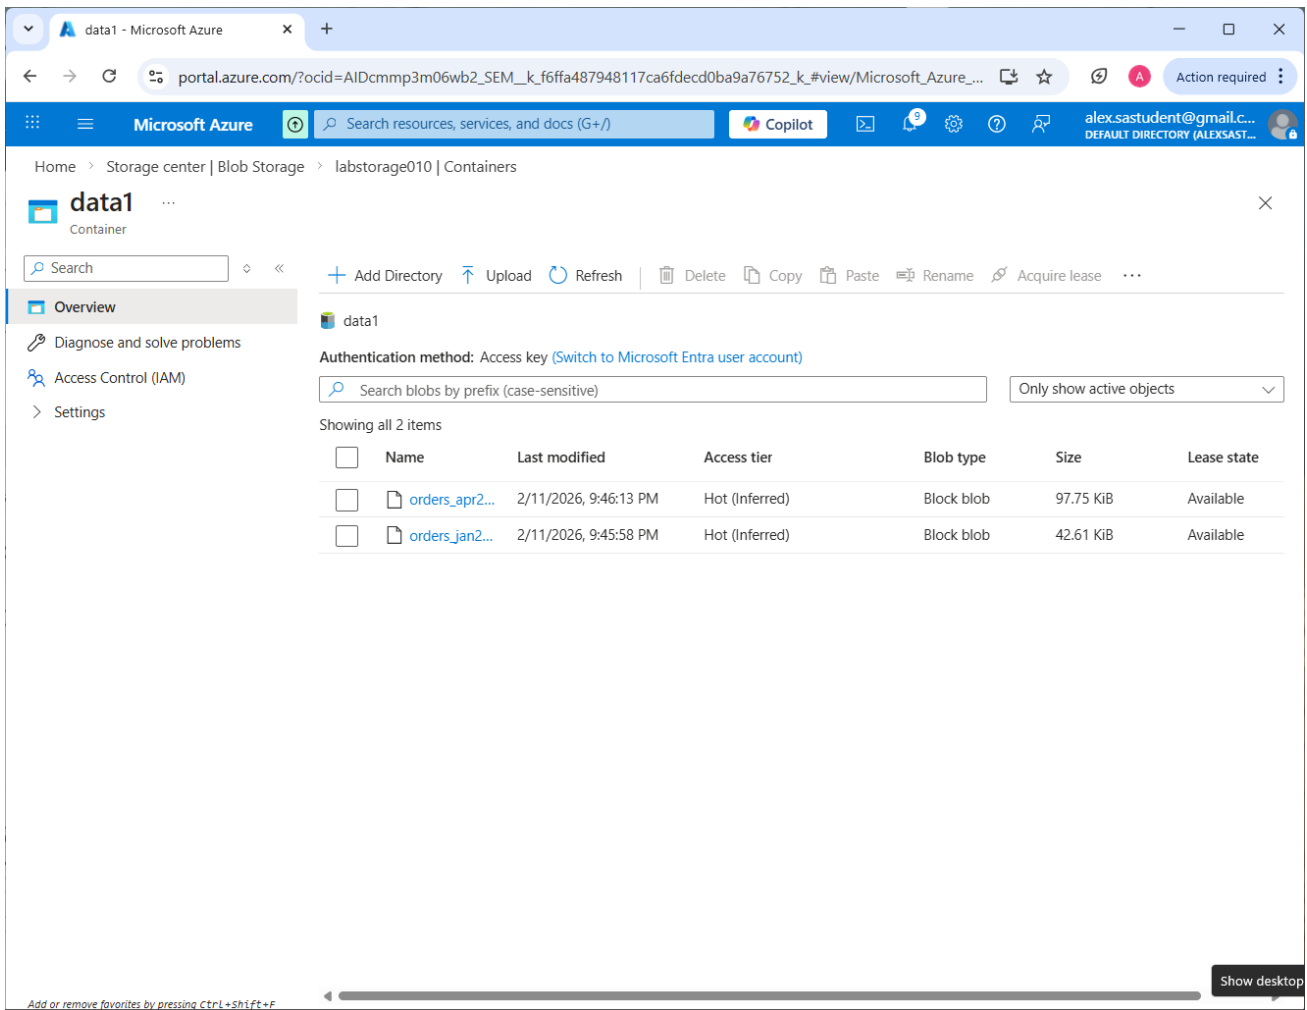
 At the bottom left, there is a small note: "Add or remove favorites by pressing Ctrl+Shift+F".

The screenshot displays the Microsoft Azure portal interface. On the left, a navigation pane shows the 'Containers' section selected under the 'labstorage010' storage account. The main area shows a list of containers with one item, '\$logs', listed. Overlaid on the right is the 'New container' dialog box. The 'Name' field is filled with 'data'. The 'Anonymous access level' dropdown is set to 'Private (no anonymous access)'. A blue information box below the dropdown states: 'The access level is set to private because anonymous access is disabled on this storage account.' At the bottom of the dialog, there is a blue 'Create' button and a 'Give feedback' link.



## Activities 3.2: Upload CSV files from Azure Portal

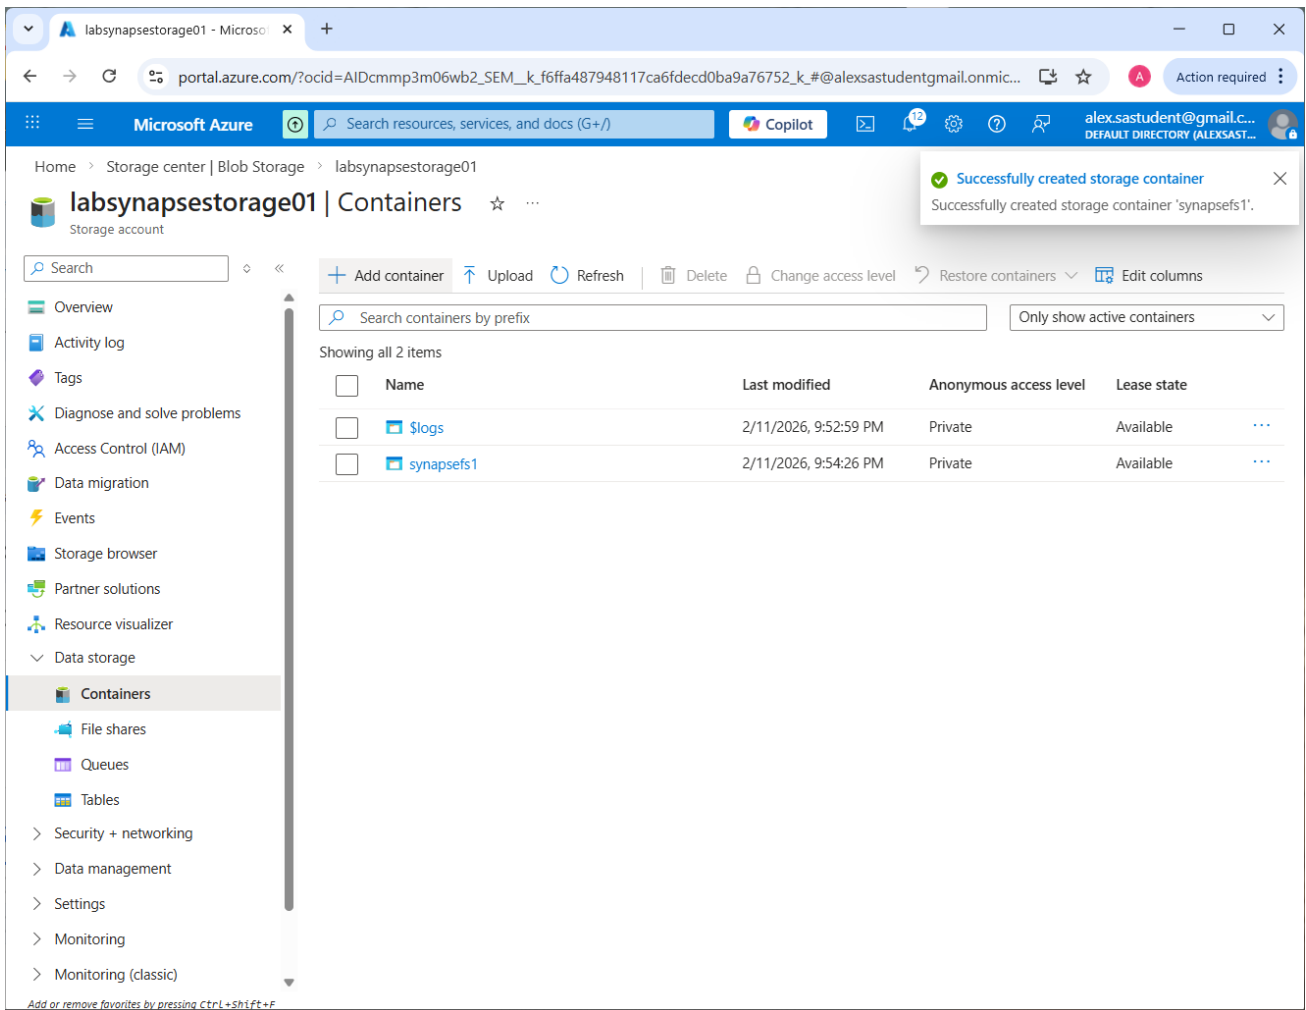
The CSV files were uploaded to an Azure Blob container.



### Activities 3.3: Ingest Data into Synapse

A Linked Service was configured in Synapse to connect to the storage account, and the Copy Data tool was used to ingest the data into a SQL table.

The screenshot shows the Microsoft Azure portal interface. At the top, the browser address bar displays the URL: `portal.azure.com/?ocid=AIDcmm3m06wb2_SEM_k_f6ffa487948117ca6fdec0ba9a76752_k_#view/HubsExtension/Deploy...`. The page title is "labsynapsestorage01\_177079933397 | Overview". The main content area features a green checkmark icon and the text "Your deployment is complete". Below this, deployment details are listed: "Deployment name: labsynapsestorage0...", "Subscription: Azure subscription 1", and "Resource group: rg-lab". The start time is "2/11/2026, 9:52:27 PM" and the correlation ID is "32652a41-f7e0-44cc-9d85-66c...". A "Go to resource" button is visible. On the right sidebar, there are sections for "Cost Management", "Microsoft Defender for Cloud", "Free Microsoft tutorials", and "Work with an expert".



### Activities 3.4: Run SQL Queries in Synapse

SQL queries similar to those used in Athena were executed in Synapse to validate the data migration. The output confirmed successful ingestion.

The screenshot shows the Azure Synapse Studio interface. The top navigation bar includes 'Microsoft Azure | Synapse Analytics | synapse-ws100'. The main workspace is titled 'Develop' and contains a file named 'SQL script 1'. The script content is as follows:

```

1 CREATE DATABASE analyticsdb1;
2 GO
3
4 USE analyticsdb1;
5 GO
    
```

The 'Properties' pane on the right shows the following details:

- Name:** SQL script 1
- Description:** (empty)
- Type:** .sql script
- Size:** 0 bytes
- Results settings per query:** First 5000 rows (default)

The 'Results' pane at the bottom displays a message: 'No results to show. Your query yielded no displayable results.' A status bar at the bottom indicates '00:00:02 Query executed successfully.'

The screenshot shows the Azure Synapse Studio interface. The top navigation bar includes 'Microsoft Azure | Synapse Analytics | synapse-ws100'. The main workspace is titled 'Develop' and contains a file named 'SQL script 1'. The script content is as follows:

```

1 CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'StrongP@ssw0rd!2026';
2 GO
    
```

The 'Properties' pane on the right shows the following details:

- Name:** SQL script 1
- Description:** (empty)
- Type:** .sql script
- Size:** 0 bytes
- Results settings per query:** First 5000 rows (default)

The 'Results' pane at the bottom displays a message: 'No results to show. Your query yielded no displayable results.' A status bar at the bottom indicates '00:00:00 Query executed successfully.'

The screenshot shows the Microsoft Azure Synapse Analytics 'Develop' environment. The SQL editor contains the following code:

```

1 CREATE DATABASE SCOPED CREDENTIAL cred_blob_sas
2 WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
3 SECRET = 'sp=r&st=2026-02-11T09:50:24Z&se=2026-02-11T18:05:24Z&sv=2024-11-04&sr=b&sig=6xpEC1ke180Fvz4FdP8em3nYm4XH51ePbtJyPma7Xs...'
4 GO

```

The 'Results' pane displays a message: "No results to show. Your query yielded no displayable results." The status bar at the bottom indicates "00:00:00 Query executed successfully."

The screenshot shows the Microsoft Azure Synapse Analytics 'Develop' environment. The SQL editor contains the following code:

```

1 CREATE EXTERNAL DATA SOURCE eds_orders
2 WITH (
3     LOCATION = 'https://labstorage0100.blob.core.windows.net/data',
4     CREDENTIAL = cred_blob_sas
5 );
6 GO
7

```

The 'Results' pane displays a message: "No results to show. Your query yielded no displayable results." The status bar at the bottom indicates "00:00:00 Query executed successfully."

The screenshot shows the Microsoft Azure Synapse Analytics 'Develop' environment. A SQL script named 'SQL script 1' is being executed. The query is a SELECT statement with a TOP 20 clause, pulling data from an OPENROWSET function. The results are displayed in a table view below the query editor.

```

1 SELECT TOP 20
2   Category,
3   Sales
4 FROM OPENROWSET(
5   BULK 'orders_jan2017.csv',
6   DATA_SOURCE = 'eds_orders',
7   FORMAT = 'CSV',
8   PARSER_VERSION = '2.0',
9   FIRSTROW = 2
10 )
11 WITH (
12   Category VARCHAR(100) COLLATE Latin1_General_100_CI_AS_SC_UTF8,
13   Sales VARCHAR(50)
14 ) AS o;
15

```

**Results**

Category	Sales
440	CA-2017-157252
516	CA-2017-127432
517	CA-2017-127432
518	CA-2017-127432

00:00:00 Query executed successfully.

## Task 4: Final Report on AWS Services and Architecture

### Activities 4.1: Report writing

#### Introduction

In this assessment, an automated cloud-based data pipeline was designed using AWS services to simulate a real-world data engineering workflow. The main objective was to ingest, catalogue, query, visualise, and monitor datasets using cloud-native services instead of traditional manual methods. The solution also included exporting curated results to Microsoft Azure Synapse Analytics to demonstrate multi-cloud interoperability.

This task helped me understand how AWS services can be integrated together to build a scalable and automated data infrastructure using service-based tools. Instead of writing large amounts of custom code, the system was built using managed services such as IAM, S3, Glue, Athena, QuickSight, Lambda, and SNS. This directly aligns with Learning Outcome 3 (LO3), which focuses on designing automated IT infrastructure using user-friendly programming and cloud services.

#### Overview of AWS Services Used

##### *AWS Identity and Access Management (IAM)*

IAM was used to control access to AWS resources securely. An IAM user was created and assigned appropriate permissions to manage S3, Glue, Athena, and other services. This step ensured that all actions performed within the environment were authorised and traceable.

Using IAM helped me understand the importance of role-based access control in cloud environments. Without proper permissions, services such as Glue and Athena cannot interact with S3 correctly.

##### *Amazon S3 (Simple Storage Service)*

Amazon S3 acted as the main storage layer of the data pipeline. A bucket was created to store the CSV datasets, and folders were organised using partition-style naming (for example, `snapshot_day=2017-01-01`).

S3 served as the raw data layer, where files were stored in their original format before being processed. Structuring folders properly was important because Glue used these paths to detect partitions automatically.

##### *AWS Glue*

AWS Glue was used for automated data cataloguing. A Glue database (`db_yoobee`) and a crawler (`yoobee_crawler`) were configured to scan the S3 bucket and detect schema automatically.

When a new folder was added (for example, `snapshot_day=2017-01-04`), the crawler was re-run, and new partitions were created without manual schema updates. This demonstrated how automation reduces repetitive work in data engineering.

### Amazon Athena

Athena was used as a serverless query engine to run SQL queries directly on data stored in S3. It connected to the Glue Data Catalog to read table metadata.

Queries such as summing sales by category and filtering by snapshot date were executed without creating a database server. Athena follows a pay-per-query model, which makes it cost-efficient for analytics tasks.

### Amazon QuickSight

Amazon QuickSight was used for data visualisation. After connecting it to Athena and authorising S3 access, dashboards were created to visualise sales by category and profit by city.

QuickSight uses SPICE to import datasets and generate visual insights without writing additional code. This demonstrates how business intelligence tools can directly integrate with cloud analytics services.

## How These Services Work Together (End-to-End Workflow)

The process begins when CSV files are uploaded into the S3 bucket. This upload triggers the Lambda function automatically. If the file extension is valid, SNS sends a notification. Meanwhile, the Glue crawler scans the S3 folder structure and updates the table metadata. Athena then queries the updated table, and QuickSight connects to Athena for visualisation. Finally, selected outputs are transferred to Azure Synapse for additional analysis.

## Automation and No-Code Features

One of the most important aspects of this project was automation without heavy coding.

The Glue crawler automatically detected schema and partitions without requiring manual table creation. Athena allowed SQL queries to be executed directly on S3 data without provisioning a database server. QuickSight allowed me to create dashboards using a graphical interface instead of writing additional code.

This approach reduced manual configuration and eliminated the need to provision virtual machines or manage infrastructure. The system mainly relied on managed services provided by AWS.

Through this process, I understood how modern cloud platforms reduce complexity by offering service-based automation rather than traditional server management.

## Data Engineering Perspective

From a data engineering perspective, the solution followed a layered approach. Raw data was stored in S3, metadata was managed through Glue, and queries were executed using Athena. Visual insights were generated in QuickSight, while Lambda and SNS handled event-driven monitoring. Seeing these components connected helped me understand how real-world data pipelines are structured.

Partitioned data design improved query efficiency by reducing the data scanned size. This is important in real-world environments where cost optimisation matters.

The pipeline also supports scalability. If new data snapshots are uploaded, the crawler updates partitions automatically, and queries continue to work without redesigning the system.

This is similar to how enterprise data pipelines are designed in real-world environments.

Before this task, I understood these concepts theoretically, but implementing them in a connected workflow made the architecture clearer to me.

## Athena vs Azure Synapse SQL – Brief Comparison

During this assessment, both Amazon Athena and Azure Synapse SQL were used for querying data. Athena is fully serverless and requires minimal setup. It directly queries data stored in S3 using the Glue catalog. It is simple and fast for analytics tasks.

Azure Synapse SQL also supports serverless querying, but requires linking storage accounts and creating external tables. It offers stronger integration within the Azure ecosystem.

From my experience, Athena was easier to configure and quicker to use for simple analytical queries. Synapse required additional configuration, such as linking storage accounts and creating external tables. However, Synapse felt more structured for enterprise-level environments. Both platforms operate on a pay-per-query model and support large-scale analytics without managing servers.

## Challenges and Solutions

Several challenges were encountered during implementation.

One issue was related to IAM permissions. Initially, Athena could not access the S3 bucket due to missing permissions. This was resolved by attaching the correct policies to the IAM role used by Glue and Athena.

Another issue occurred when the Glue crawler did not detect new partitions after adding a new folder. This happened because the folder structure was not consistent with partition naming. After correcting the naming format (snapshot\_day=YYYY-MM-DD), the crawler updated successfully. QuickSight initially showed S3 access errors. This was resolved by authorising QuickSight to access the S3 bucket in security settings.

Additionally, Athena returned empty results when querying a newly added dataset. After re-running the crawler and refreshing metadata, the issue was fixed.

These challenges improved my understanding of cloud permissions, schema management, and service integration.

At one point, I assumed the table had updated correctly because the crawler completed successfully. However, when I ran the Athena query, the results did not reflect the new data. After reviewing the partition structure, I realised the folder naming format was slightly inconsistent. Correcting the folder naming and re-running the crawler resolved the issue.

## Summary and Learning Outcome Reflection

Overall, this project gave me practical experience in building a serverless data pipeline using AWS services. Each service played a specific role, and together they formed a scalable and cost-effective architecture.

Through this implementation, I gained practical experience in configuring IAM, designing S3 storage structures, automating schema detection with Glue, querying with Athena, and visualising data using QuickSight. The addition of Lambda and SNS further demonstrated event-driven automation. This project helped me to be familiar with designing and automating infrastructure using cloud-based services rather than manual configuration. It reflects real-world cloud engineering practices and strengthened my understanding of multi-cloud data architecture.